ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct1310

Artificial Intelligence in Cloud Security: Techniques, Challenges, and Future Directions

Mohamed Riyaz M. Meera Rawuthar; Ali M. Iqbal

Enterprise Digital Solutions Division, Saudi Aramco, Dhahran, Saudi Arabia

Publication Date: 2025/10/31

Abstract: Cloud platforms deliver elasticity and scale but also widen the attack surface via multi-tenancy, rapid change, and opaque dependencies. This narrative survey synthesizes peer-reviewed work (2020–2025) and major-vendor documentation on how artificial intelligence (AI)—including anomaly detection, intrusion detection systems (IDS), user and entity behavior analytics (UEBA), privacy-preserving/federated learning (FL), and reinforcement learning (RL)—strengthens cloud defense. Evidence across recent studies indicates that (i) supervised and unsupervised learning detect previously unseen behaviors beyond signature baselines; (ii) Shapley-value explanations for log anomalies can improve analyst triage with minimal accuracy loss; (iii) FL with secure/verifiable aggregation and differential privacy reduces raw-data exposure but remains vulnerable to poisoning and Byzantine behaviors; and (iv) RL can automate containment/response steps in closed-loop SOC workflows. Persistent challenges include dataset shift and class imbalance, adversarial robustness, and latency/cost at cloud scale. We outline directions in robust/verified FL, lightweight edge–cloud models, graph learning for threat intelligence, and standardized cloud-native benchmarks with calibration and latency reporting. No new experiments were conducted; we provide a structured synthesis and an explicit selection protocol.

Keywords: *Cloud Security; Intrusion Detection; Anomaly Detection; User & Entity Behavior Analytics (UEBA); Explainable AI; Federated Learning; Differential Privacy; Reinforcement Learning.*

How to Cite: Mohamed Riyaz M. Meera Rawuthar; Ali M. Iqbal (2025) Artificial Intelligence in Cloud Security: Techniques, Challenges, and Future Directions. *International Journal of Innovative Science and Research Technology*, 10(10), 1880-1883. https://doi.org/10.38124/ijisrt/25oct1310

I. INTRODUCTION

Public cloud services underpin modern digital ecosystems by enabling rapid scaling, faster delivery cycles, and usage-based cost models. The same elasticity and multi-tenancy expand the attack surface, exposing tenants to data breaches, insider abuse, and agile APTs that evolve faster than static controls. Traditional defenses—perimeter firewalls, rule-based IDS, and manual monitoring—struggle with dynamic workloads and noisy, heterogeneous telemetry.

AI methods offer a complementary approach: learning behavioral baselines across identity, network, process, and control-plane logs; surfacing weak signals amid noise; and automating parts of investigation and response. When integrated with cloud-native architectures, these capabilities can shift security operations from reactive to proactive while reducing mean time to detect/respond. At the same time, AI introduces risks, including adversarial manipulation, privacy leakage when training on sensitive logs, and reduced transparency in model decisions. This survey reviews recent AI techniques for cloud security, evaluates operational readiness, and identifies open challenges and near-term opportunities.

II. LITERATURE REVIEW

Machine-learning-based anomaly and intrusion detection in cloud settings shows gains on public corpora and real logs. Representative studies include Parameswarappa et al. on cloud-oriented anomaly detection [1] and Attou et al. on ML-based cloud IDS [2]. Beyond accuracy, explainability improves triage and analyst trust; Alam et al. (SXAD) demonstrate Shapley-style explanations for log anomaly detection with maintained performance (note: the paper title uses “Shapely”) [3]. RL is emerging to automate elements of incident-response playbooks and penetration-testing workflows [4]. Privacy-preserving analytics combine federated learning (FL) with secure/verifiable aggregation and differential privacy; surveys systematize attacks (poisoning, Byzantine, inference) and defenses [5], [7], while system papers examine robustness under secure aggregation and verifiable aggregation protocols [8], [9]. Broader surveys on AI/ML in cybersecurity provide context on method families and deployment factors [6].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct1310

III. REVIEW METHODOLOGY

- *Scope*

We covered 2020–2025 peer-reviewed publications (IEEE, ACM, Springer) and official cloud-provider documentation (Google Cloud, Microsoft, Alibaba).

- *Databases*

IEEE Xplore, ACM Digital Library, SpringerLink, Google Scholar.

- *Query Blocks*

("cloud" OR "cloud security" OR "cloud threat") AND (anomaly detection OR IDS OR UEBA OR "behavior analytics" OR "federated learning" OR "secure aggregation" OR "reinforcement learning" OR XAI).

- *Inclusion*

(i) Cloud-relevant telemetry (VPC/network, IAM/auth, VM/container, control plane); (ii) AI/ML methods for detection/triage/response; (iii) FL security surveys/systems; (iv) major-vendor docs for managed detections/UEBA [12]–[16].

- *Exclusion*

(i) Non-archival preprints unless documenting datasets/deployed systems; (ii) purely on-prem IDS; (iii) opinion pieces.

- *Data Charting*

For methods papers we extracted use case, data modality, model family, evaluation setup (dataset, temporal split), metrics (including calibration where available), and latency/cost notes. For provider docs we recorded detector classes and intended latency windows.

- *Synthesis*

Narrative (no meta-analysis).

- *Limitations*

Narrative scope; no new experiments; constrained by limited cloud-native, openly reported latency/cost metrics.

IV. CLOUD SECURITY FUNDAMENTALS

Cloud security protects confidentiality, integrity, and availability across distributed, multi-tenant architectures. Prominent risks include misconfiguration, insecure APIs, weak identity and access management, insider abuse, and data exfiltration. Providers and customers share responsibility: the provider secures the cloud infrastructure; customers secure workloads, identities, and data.

A. Threat Model and Scope

We consider attackers ranging from opportunistic actors (credential-stuffing, cryptomining) to APTs with cloud-lateral-movement capabilities. We scope to threats observable in cloud telemetry (e.g., IAM anomalies, VM/process behavior, control-plane/API misuse, network-flow deviations) and to responses achievable with cloud-native controls (isolation, policy changes, key rotation).

B. Why AI in Cloud?

Static signature- or rule-based controls struggle with novel behaviors and degrade under drift. AI augments existing controls by learning baselines, detecting deviations in (near) real time, and automating elements of investigation and containment.

V. LEVERAGING AI FOR INTELLIGENT CLOUD DEFENSE

AI contributes pattern recognition over heterogeneous telemetry (network, identity, application/process, control-plane logs) and decision-making for containment. Supervised models classify known threats; unsupervised and self-supervised representations surface unknown anomalies; RL can recommend or execute mitigations subject to guardrails. Embedded into SOC workflows, these capabilities support a shift from playbook-driven response to closed-loop detection and response. Where latency and cost constraints apply, stream processing and model compression become first-class design choices.

VI. AI TECHNIQUES FOR CLOUD SECURITY

A. Anomaly Detection

Autoencoders, variational autoencoders (VAEs), and isolation-style methods learn baselines over high-dimensional logs and surface deviations. Pairing detections with local explanations (e.g., Shapley-value attributions over features) improves triage without materially degrading performance [1], [3].

B. Intrusion Detection Systems (IDS)

Deep models (e.g., CNNs, LSTMs, stacked autoencoders) operate on network flows and application-layer features to identify intrusion patterns that evade signatures. In cloud environments, IDS should report detection latency and cost alongside accuracy to reflect autoscaling realities [2], [3].

C. Behavioral Analytics (UEBA)

UEBA learns per-identity and per-asset behavior profiles (e.g., login times/locations, access paths, data movement) and flags anomalies, enriching incidents with context. Sequence models supplemented with graph features can reduce false positives while preserving sensitivity [3].

D. Automated Response Systems

RL and policy-learning approaches can map alerts and context to containment actions (isolate a VM, revoke a token, rotate a key) subject to constraints and human approval. Safety requires simulation/sandboxing and conservative policies for production [4].

E. Privacy-Preserving Learning

Federated learning (FL) with secure aggregation and differential privacy enables cross-tenant or cross-region learning without centralizing raw data. Robustness against poisoning/Byzantine clients and verifiability of aggregation remain active challenges [5], [7]–[9].

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct1310

F. Graph Learning for Threat Intelligence

Graph neural networks (GNNs) can encode relationships among identities, assets, IPs, processes, and indicators of compromise. By scoring nodes/edges or subgraphs, GNNs support correlation and prioritization across disparate signals—useful for lateral-movement and multi-stage campaigns.

VII. CHALLENGES AND LIMITATIONS

A. Data Quality, Shift, and Labeling

Cloud-native attacks are rare and evolve quickly; popular corpora (e.g., NSL-KDD, BoT-IoT) are imbalanced and dated for cloud traffic if used naïvely. Prefer temporal splits, cost-sensitive metrics, calibration (e.g., ECE/Brier), and ablations instead of headline accuracy alone. **[10]**

B. Adversarial Vulnerabilities

Models can be evaded or poisoned by inputs designed to induce false negatives/positives. Defenses include adversarial training, input validation, uncertainty-aware scoring, and ensembling; none are foolproof under open-world cloud telemetry. (See also FL robustness in VIII-E.)

C. Computational Cost and Latency

Deep models can be costly in autoscaled environments. Apply model compression/distillation and streaming inference to meet near-real-time SLAs and budget constraints.

D. Privacy and Ethics

Training on sensitive logs risks privacy leakage or bias amplification. FL with secure aggregation and differential privacy mitigates raw-data exposure but needs careful threat modeling and utility–privacy trade-off calibration. [5]

E. Explainability and Trust

Domain-tuned explanations (e.g., SHAP on IAM anomalies) and human-in-the-loop review can increase analyst trust with minimal accuracy loss. (Use "Shapley-value explanations" in text; keep source title spelling in the reference.)

VIII. CASE STUDIES / REAL-WORLD APPLICATIONS

A. Google Cloud Security Command Center (SCC)

Event Threat Detection (ETD) provides continuously updated, log-based detectors surfaced as findings in near real time (Premium tier). Virtual Machine Threat Detection (VMTD) adds agentless memory scanning to flag cryptomining malware and kernel-mode rootkits in Compute Engine VMs. Public documentation indexes detector coverage and response guidance [12], [13].

B. Microsoft Sentinel (SIEM/SOAR) with UEBA

Sentinel includes UEBA that builds behavioral baselines for identities/devices and enriches incidents with anomalies. Configuration and operations are exposed through Defender/Sentinel portals and documented reference/enablement guides [14], [15].

C. Alibaba Cloud Security Center (CNAPP) and PAI

Security Center consolidates asset/posture risk detection and threat alerts (e.g., suspicious processes, unusual logons, webshells) using detection models and log analytics. PAI offers ML anomaly-detection components for customer pipelines; public docs emphasize CSPM/CWPP and ML-driven checks [16].

Implication. Major providers combine rule-based signals with ML/behavior analytics and managed detections. For comparability, researchers should document detector classes and latency ceilings when claiming "real-time."

IX. FUTURE DIRECTIONS AND RESEARCH OPPORTUNITIES

A. Adversarial Robustness

Advance attack/defense evaluations under cloud-realistic constraints (streaming, partial observability). Integrate adversarial training, input sanitization, and uncertainty-aware scoring into cloud-native pipelines.

B. Explainable AI for SecOps

Design domain-specific explainers and UX that improve triage speed and reduce alert fatigue, with human-in-the-loop validation.

C. Federated and Privacy-Preserving Learning

Pursue robust/verified FL: secure/verifiable aggregation, poisoning-resilient optimization, and DP configurations that preserve utility for rare classes, especially across multi-cloud/edge.

D. AI-Augmented Threat Intelligence

Use graph learning to correlate indicators, vulnerabilities, and behaviors across identities and assets, improving prioritization and lateral-movement detection.

E. Lightweight Models for Edge–Cloud

Adopt model compression/distillation and streaming inference to meet latency and cost budgets while maintaining detection quality.

F. Standardization and Benchmarking

Develop cloud-native benchmarks with open telemetry, evaluation protocols, and reporting standards (latency, cost, calibration), reducing reliance on dated corpora.

X. CONCLUSION

AI measurably advances cloud defense across anomaly detection, IDS, UEBA, and privacy-preserving learning. The main blockers are dataset shift/imbalance, adversarial brittleness, opaque decision-making, and cost/latency at scale. Near-term wins include verified/robust FL (e.g., secure input checks plus verifiable aggregation), domain-tuned explainability for logs and identities, and lightweight models at the edge–cloud boundary. Standardized cloud-native benchmarks—and routine reporting of latency, cost, and calibration—will help translate research findings into deployment-ready practice.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25oct1310

REFERENCES

[1]. Parameswarappa, P.; Shah, T.; Lanke, G.R. "A Machine Learning-Based Approach for Anomaly Detection for Secure Cloud Computing Environments." Proc. 2023 Int. Conf. on Intelligent Data Communication Technologies and Internet of Things (IDCIoT), IEEE, 2023. https://doi.org/10.1109/IDCIoT56793.2023.10053518

[2]. Attou, H.; Guezzaz, A.; Benkirane, S.; Azrour, M.; Farhaoui, Y. "Cloud-Based Intrusion Detection Approach Using Machine Learning Techniques." Big Data Mining and Analytics 6(3), 311–320, 2023. https://doi.org/10.26599/BDMA.2022.9020038

[3]. Alam, K.; Kifayat, K.; Sampedro, G.A.; Karovič, V.; Naeem, T. "SXAD: Shapely eXplainable AI-Based Anomaly Detection Using Log Data." IEEE Access 12, 95659–95672, 2024. https://doi.org/10.1109/ACCESS.2024.3425472

[4]. Rahman, A.; Redino, C.; Nandakumar, D.; Cody, T.; Shetty, S.; Radke, D. Reinforcement Learning for Cyber Operations: Applications of Artificial Intelligence for Penetration Testing. Wiley–IEEE Press, 2025. https://doi.org/10.1002/9781394206483

[5]. Cunha Neto, H.N.; Hribar, J.; Dusparic, I.; Mattos, D.M.F.; Fernandes, N.C. "A Survey on Securing Federated Learning: Analysis of Applications, Attacks, Challenges, and Trends." IEEE Access 11, 41928–41953, 2023. https://doi.org/10.1109/ACCESS.2023.3269980

[6]. Ozkan-Okay, M.; Akin, E.; Aslan, Ö.; Kosunalp, S.; Iliev, T.; Stoyanov, I.; Beloev, I. "A Comprehensive Survey: Evaluating the Efficiency of Artificial Intelligence and Machine Learning Techniques on Cyber Security Solutions." IEEE Access 12, 12229–12256, 2024. https://doi.org/10.1109/ACCESS.2024.3355547

[7]. Hu, K. et al. "An overview of implementing security and privacy in federated learning." Artificial Intelligence Review (2024). https://doi.org/10.1007/s10462-024-10846-8

[8]. Lycklama, H.; Burkhalter, L.; Viand, A.; Küchler, N.; Hithnawi, A. "RoFL: Robustness of Secure Federated Learning." Proc. IEEE Symposium on Security and Privacy (S&P), 2023, pp. 453–476. https://doi.org/10.1109/SP46215.2023.10179400

[9]. Eltaras, T.; Sabry, F.; Labda, W.; Alzoubi, K.; Malluhi, Q. "Efficient Verifiable Protocol for Privacy-Preserving Aggregation in Federated Learning." IEEE Transactions on Information Forensics and Security 18, 2977–2990, 2023. https://doi.org/10.1109/TIFS.2023.3273914

[10]. Tavallaee, M.; Bagheri, E.; Lu, W.; Ghorbani, A.A. "A Detailed Analysis of the KDD CUP 99 Data Set." IEEE CISDA, 2009. https://doi.org/10.1109/CISDA.2009.5356528

[11]. Koroniotis, N.; Moustafa, N.; Sitnikova, E.; Turnbull, B. "Towards the Development of Realistic Botnet Dataset in the Internet of Things for Network Forensic Analytics: BoT-IoT Dataset." arXiv:1811.00701, 2018. https://arxiv.org/abs/1811.00701

[12]. Google Cloud. "Overview of Event Threat Detection (ETD) — Security Command Center." 2025. https://cloud.google.com/security-command-center/docs/concepts-event-threat-detection-overview

[13]. Google Cloud. "Virtual Machine Threat Detection (VMTD) overview — Security Command Center." 2025. https://cloud.google.com/security-command-center/docs/concepts-vm-threat-detection-overview

[14]. Microsoft Learn. "Microsoft Sentinel — User and Entity Behavior Analytics (UEBA) reference." 2025. https://learn.microsoft.com/en-us/azure/sentinel/ueba-reference

[15]. Microsoft Learn. "Enable User and Entity Behavior Analytics (UEBA) in Microsoft Sentinel." 2025. https://learn.microsoft.com/en-us/azure/sentinel/enable-entity-behavior-analytics

[16]. Alibaba Cloud. "What is Security Center (CNAPP)." 2025. https://www.alibabacloud.com/help/en/security-center/product-overview/what-is-security-center